RAID - A Risky Standard for Video Storage

RAID has become a risky standard for video storage. RAID maximizes the stresses on hard disks (temperature, vibration & wear), thus producing a high probability of disk failure, even with more expensive drives. With ever-increasing storage capacities, the rates of error increase, as do RAID rebuild times. In fact, RAID rebuild times become so extended with the highest capacity drives that there is a significant probability of a total and catastrophic loss of all data in the array. System integrators are learning this the hard way as higher and higher array capacities are needed for modern high-definition surveillance systems. Today, there is a new technology that directly addresses these problems, offering a much more reliable storage solution with many advanced features and zero risk of total data loss. The name of this solution is COLDSTORE.

RAID Risks

- Reduced HDD lifespan due to constant reading and writing, heat, and vibrations
- Performance downgrades during rebuild period
- Dual HDD failure results in total loss of data
- Hot Spares are activated with reduced lifetime
- Vigilant monitoring of HDD status is required to mitigate the possibility of dual HDD failure

The Risks of RAID-5 and RAID-6

RAID modes 5 and 6 distribute data across multiple hard disk drives (HDD). To distribute the data the RAID controller must constantly write to each HDD in a complex process, meaning the entire array of disks is working at maximum duty cycles. Constant HDD activity means increased temperature, high levels of vibration and increased wear, resulting in reduced disk lifespan.

RAID's greatest benefit is that it offers redundancy. Whenever a drive fails in a RAID-5 array the failed HDD can be rebuilt using the data stored on other HDDs. This method of creating redundancy comes at a price. During a rebuild the server will experience noticeable decline in performance. HDD capacity has grown at a much faster rate than write speed. Larger capacity drives now take many hours - if not days - to rebuild, meaning long periods of degraded performance.

During the rebuild period the biggest worry is not the decline in system performance, though this should be considered, but that a second disk will fail during the same period. Should a second drive fail while the first is still being rebuilt, the data stored in the RAID array will be unrecoverable. Many RAID-5 arrays with high disk capacities are in almost constant rebuild, with a high risk of total data loss if a second fails during this process. RAID-6 copes with up to two drives failing, with a much lower read & write performance (especially during rebuild), at the expense of processing power, complexity and system cost.

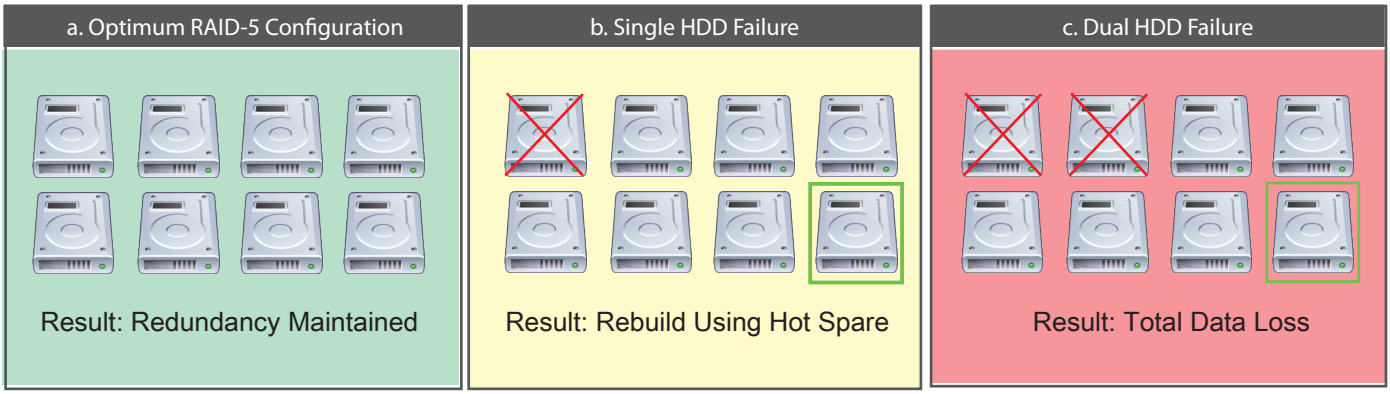

Figure 1.1 RAID-5 Configuration with Single and Dual HDD Failure (a) All HDDs are functioning and the RAID configuration retains redundancy, but at maximum HDD stress. (b) A single HDD fails and the hot spare must become active (green square) and the array undergoes a rebuild. The data is still accessible; however, the recorder's resources will be allocated to the rebuild leaving less computing resources available for recording and playing video. If another HDD fails during the rebuild, then the recorder will lose all data. (c) In case of dual HDD failure, all the data on the array will be lost.

RAID servers usually include so-called "hot spares", which is a stand-by extra disk ready to take over if a disk fails. This means that in the event of a disk error, a rebuild may begin automatically in the "background". Theoretically, the hot spare will be "new" but in fact it will have been heated and vibrated in a very similar pattern to the rest of the HDD in the array and so in many ways it is just as likely to fail as any other "used" HDD in the array. Because RAID-5/6 can undergo automatic rebuild, the system's status needs to be constantly monitored as the risk of total data loss is very high. In such cases all the disks may have to be replaced to ensure more reliable operation. This requires on-site intervention by a trained operator.

Coldstore Benefits

- Eliminates the risk of dual HDD failure associated with RAID-5
- Reduces data reading/writing by 87% to prolong the lifespan of your HDDs
- Requires no rebuild time in the event of HDD failure providing continuous high performance
- Runs over 30TB of storage on a measly 50 W on standard SATA disks
- Records sequentially for easy data retrieval

It's as Easy as 1, 2, 3 – LAID, SFS, and Data Mirroring

COLDSTORE exploits a revolutionary patented storage system technology incorporating a Linear Array of Idle Disks (LAID™), a special Sequential Filing System (SFS™) to store data sequentially within the disk and across disks, and a unique overlapping mirrored-pair writing pattern. This complete system has been specifically designed to solve the problems of high-resolution video surveillance storage.

SFS™ writes data sequentially across the entire disk, exploiting the sequential nature of video recording. SFS™ reduces HDD operating temperature, eliminates vibration and reduces wear to an absolute minimum, thus increasing the reliability of individual HDDs.

LAID™ writes across the disk array sequentially, switching off all disks not being written to. This dramatically reduces disk wear and duty cycle, and reduces the total power required by the array (and therefore reduces the operating temperatures and the heat generated).

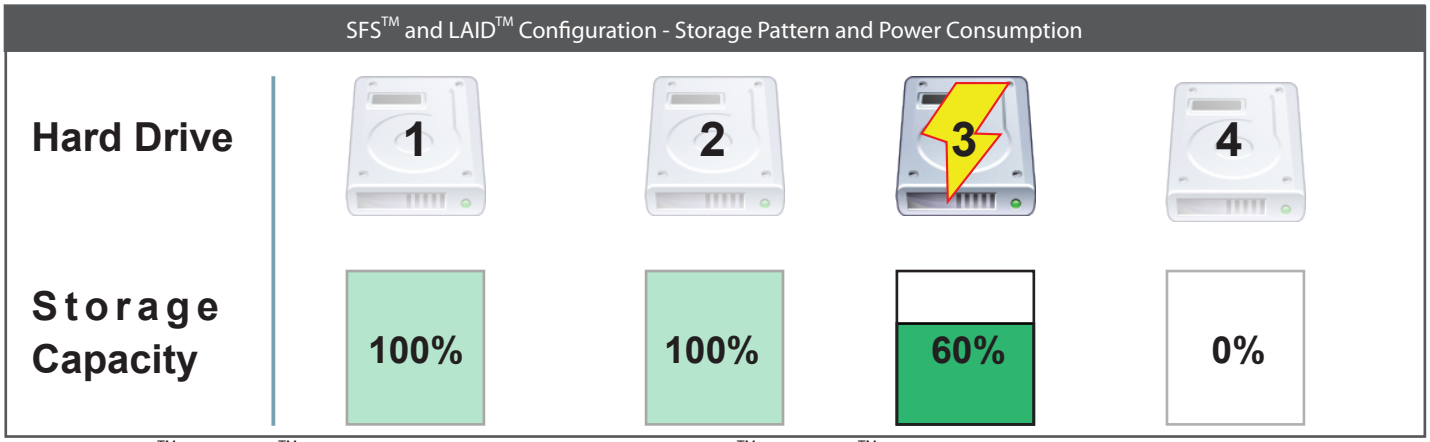

Figure 1.2 SFS™ and LAID™ HDD Configuration - In this situational SFS™ and LAID™ configuration, HDDs 1, 2, and 4 are not being actively storing or retrieving data, so they have been switched off to reduce power consumption and extend the lifetime of the HDDs. HDD 1 and 2 have filled capacity while HDD 4 remains empty. HDD 3 is actively being written to, therefore has power. Once HDD 3 is filled to capacity it will turn off, just like HDD 1 and 2, and HDD 4 will power on and remain active until filled to capacity.

The overlapping mirrored-pair writing pattern writes identical data to two consecutive disks, providing complete data redundancy during the critical writing process. Once a mirrored-pair is filled, the first disk is switched off, safely retaining the data. The second remains on and a third disk is switched on and the mirrored writing position shifts one disk position along, now writing to the second and third disks, which are filled in turn. The process then repeats, sequentially stepping along the disks in the array. Should a disk fail, the data is preserved on the good disk of the pair and the writing point simply steps on to the next pair of disks earlier than normal.

This overlapping-pair writing process results in mirrored data writing resilience, without having to use twice the number of disks, thus reducing cost. When a disk fails, no rebuild is required as the data is already preserved. All disks except the current pair being written to are switched off, extending disk lifespan and greatly reducing power consumption and heat generation. In fact, in a fully-populated 15-bay COLDSTORE array, the disks are off for 87% of the time.

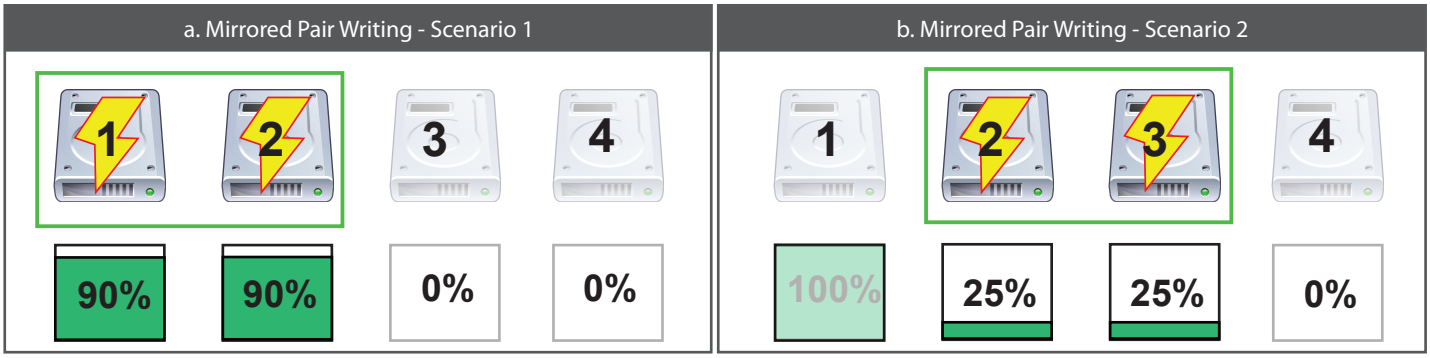

Figure 1.3 Mirrored Pair Writing - COLDSTORE Redundancy (a) HDD 1 is recording and HDD 2 is mirroring the data. In the event that HDD 1 fails during recording, HDD 2 will continue to store the data, offering redundancy. (b) HDD 1 has successfully filled to capacity. HDD 2 and 3 are continuing the data mirroring pattern. Once HDD 2 fills to capacity, HDD 3 and HDD 4 will continue the pattern.

Using these advanced techniques and specifically designed low-power, high-reliability hardware, a COLDSTORE system can offer 30TB storage capacity with a power consumption of less than 50 watts. This is less than 1.7 watts per Terabyte, and easily TEN times less than standard RAID-5 or RAID-6. In effect, COLDSTORE can provide ten times more capacity for the same power as RAID.

Environmental Friendly

The New Standard for Surveillance Data Storage

High-capacity COLDSTORE arrays, which eliminate all of the risks associated with RAID, combined with the high availability features of Instek Digital's MatriVideo surveillance recording and control platform, such as hybrid failover and power redundancy, create a powerful and extremely reliable total surveillance management solution.

INSTEK DIGITAL
www.instekdigital.com

Instek Digital USA

3661 Walnut Ave,
Chino, CA 91710, USA
TEL: (909) 627-8811
FAX 1 (909) 627-8838
Email sales@instekdigital.com

Instek Digital USA

4F, No.119-1,Baozhung Road,
Xindian City,Taipei County, Taiwan 231
TEL: +886-2-29141839
FAX: +886-2-29101803
info@instekdigital.com

V.20101110